# Formalization of Phase Ordering

Tiago Cogumbreiro
Rice University

Jun Shirako
Rice University

Vivek Sarkar
Rice University

Phasers pose an interesting synchronization mechanism that generalizes many collective synchronization patterns seen in parallel programming languages, including barriers, clocks, and point-to-point synchronization using latches or semaphores. This work characterizes scheduling constraints on phaser operations, by relating the execution state of two tasks that operate on the same phaser. We propose a formalization of Habanero phasers, May-Happen-In-Parallel, and Happens-Before relations for phaser operations, and show that these relations conform with the semantics. Our formalization and proofs are fully mechanized using the Coq proof assistant, and are available online.

## 1 Introduction

Phasers are an interesting synchronization mechanism that generalizes barriers with collective producer-consumer synchronization. A phaser can encode the synchronization mechanism of latches, futures, join barriers, cyclic barriers, as well as any *collective synchronization pattern* provided by CUDA, C♯, Java, MPI, and X10. Phasers [13] were first introduced in the Habanero Extreme Scale research project at Rice University, as an extension to X10 clocks [2], and implemented in Habanero-Java and Habanero-C. A restricted form of phasers was also introduced in the standard `java.util.concurrent.Phaser` library starting with Java 7. The phaser synchronization mechanism is relevant at the theoretical level because of its generality. Theoretical results that target phasers can easily translate across different languages and parallel runtimes [4].

The phaser synchronization mechanism lets tasks observe a collective event, called *phase*, which is visible once every member of a group of tasks *signals* the phaser exactly once. We define *signalers* of the phaser as the group of tasks able to signal a phaser. The same phaser can be used to observe multiple phases, which are distinguishable by a natural number. A task can observe phase $n$ once each signaler issues at least $n$ signals. Phaser synchronization also features *dynamic membership*, that is, the group of signalers can grow and shrink dynamically: a signaler can add a member, which in turn inherits the signal count of the task adding it; a signaler can also revoke its membership at any time.

As an example of phaser synchronization, let us consider a group of three tasks, uniquely identified by $t_1$, $t_2$, and $t_3$, and let this group of tasks be the signalers of phaser $P$. Also, lets examine a point in time, with respect to phaser $P$, where task $t_1$ signaled 3 times, task $t_2$ signaled 4 times, and task $t_3$ signaled 10 times. Tasks can use $P$ to observe any phase below or equal to phase 3, since the signalers collectively issued at least 3 signals. Conversely, at this point in time, any phase above 3 is *not* observable, *e.g.,* for phase 4 to be observed we are missing a signal from task $t_1$. Dynamic membership affects synchronization: if task $t_1$ adds a task $t_4$ as a signaler of phaser $P$, then for phase 4 to be observable we are missing a signal from task $t_1$ and a signal from task $t_4$; and, if, subsequently, tasks $t_1$ and $t_4$ revoke their membership, then phase 4 is observable.

This paper introduces the first formalization of Habanero phasers and also presents an Happens-Before (HB) [9] relation and a May-Happen-In-Parallel [6] (MHP) relation for phaser operations, both of which are fundamental problems for concurrency analysis. MHP and HB characterize scheduling restrictions between two instruction instances. An example of the HB relation is ordering any instruction

D. Orchard and N. Yoshida (Eds.): Programming Language Approaches to Concurrency- and Communication-Centric Software (PLACES 2016).
EPTCS 211, 2016, pp. 13–24, doi:10.4204/EPTCS.211.2

© T. Cogumbreiro, J. Shirako, V. Sarkar
This work is licensed under the
Creative Commons Attribution License.

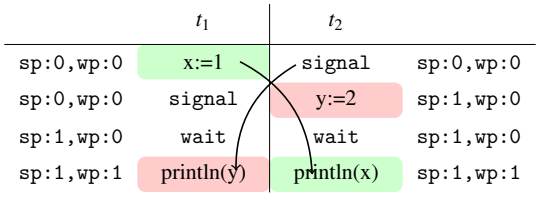


Figure 1: Phase-ordering between two task traces in a program with a race error.

that happened before spawning a task and any instruction in the task body being spawned; MHP can be defined using the HB relation. MHP and HB analysis are fundamental in the verification of barrier synchronization errors [12], lock-based deadlock prediction [1, 7], and race-detection [10, 14].

The HB relation we introduce comes from the Phase Ordering definition [13] that relates the execution state of two tasks manipulating the same phaser: if the number of signals issued by a task (property `sp`) is smaller than the last phase observed by some other task (property `wp`) then the former task happened before the latter task. The execution of a program that uses Habanero phasers must respect the scheduling restriction imposed by Phase Ordering, but how can we be sure that this property holds? The example in Figure 1 lists the execution trace of two tasks and also includes the `sp` and `wp` at each step, for both tasks, and w.r.t. the same phaser. Tasks increment their `sp` after signaling, and their `wp` after waiting. Note how HB orders instructions at different points in time: the assignment `x:=1` by $t_1$ happens before `println(x)` by $t_2$, so we can conclude that $t_2$ reads the value written by task $t_1$. Conversely, HB does not order the write `y:=2` by $t_2$ and the read `println(y)` by $t_1$, so, as read and write are unsynchronized, there is a data race. The goal of this work is twofold: 1) prove that to schedule phaser-operation across tasks it is sufficient to compare a task-local property `sp` of one task with the last global observation `wp` of another task; and 2) and prove that the semantics of Habanero phasers respects the Phase Ordering property.

Crafa *et al.* propose a Coq formalization of a subset of the X10 and define a HB relation in [5], but only consider fork-join synchronization, and omit dynamic barrier synchronization that we formalize. Tomofumi *et al.* use HB and MHP to check for data races in the polyhedral subset of clocked X10 programs [15]. Joshi *et al.* propose an informal MHP relation for X10 clocks [8].

This paper establishes two main properties with respect to the phasers semantics we introduce. First, as required by HB and MHP analysis, we show that the HB relation we define is a causality relation [9]. Second, since the HB relation is defined on the state of a phaser $P$, we show that HB *conforms* with the execution semantics of phasers; that is, if a state $P$ reduces to $Q$ after zero or more steps, then $Q$ cannot happen before $P$. By targeting phasers, our formalization unifies collective producer-consumer synchronization [12] and barriers with dynamic membership [8] in a single theoretical framework. Additionally, we formalize and establish the correctness of our definitions with proofs verified by the Coq proof assistant, available online, as part of our **HJ-Coq formalization project [3]**.

The main contributions of this paper are:

1. introduces the first formalization semantics of Habanero phasers;
2. defines an HB relation and an MHP relation for phaser operations;
3. shows that HB is a causality relation, given by Theorem 1;
4. shows that HB conforms with the reduction relation, given by Theorem 3;
5. presents the full Coq mechanization of the theory, along with examples.

In the next section, we describe the phaser operations and its semantics. In Section 3, we introduce Phase Ordering, the MHP relation, and the HB relation. Next, in Section 4, we establish the main results with regards to the semantics of phasers. We conclude in Section 5 and discuss future directions.

## 2 Phaser semantics

Let us discuss informally the semantics of Habanero phasers by revisiting the example in Figure 1. In the following Java code listing, tasks $t_1$ and $t_2$ synchronize each of their access to two different shared variables `x` and `y` by means of a phaser `ph`.

```
1  ph = newPhaser(SIG_WAIT);
2  asyncPhased(ph.inMode(SIG_WAIT), () -> {
3    ph.signal();
4    y = 2;
5    ph.doWait();
6    println(x);
7    ph.drop();
8  });
9  x = 1;
10 ph.signal();
11 ph.doWait();
12 println(y);
```

Task $t_1$ creates phaser `ph` in Line 1 and then spawns a task $t_2$ in Line 2. These two tasks are the signalers of `ph`. The creator of a task is a signaler of that phaser. Signalers, and only signalers, can register other members by spawning them with `asyncPhased` and passing the target phaser. Here, task $t_1$ registers task $t_2$ with phaser `ph` — we postpone discussing the meaning of expression `ph.inMode(SIG_WAIT)` in Line 2; for now it is enough to interpret the expression as `ph`. Task $t_1$ then writes to variable `x` in Line 9 and reads from variable `y` in Line 12, while, concurrently, task $t_2$ writes to variable `y` in Line 4 and reads from variable `x` in Line 6. Before terminating, task $t_2$ revokes its membership on phaser `ph` by invoking `ph.drop()` in Line 7.

Tasks $t_1$ and $t_2$ synchronize in the example by executing `ph.doWait()` in Lines 5 and 11: each task waits for phase 1 to be observed, which can only happen once both task execute `ph.signal()` in Lines 3 and 10. The reason there is a data race in Lines 4 and 12 is because task $t_1$, that blocks with `ph.doWait()` before reading `y`, can unblock and read the variable when task $t_2$ signals in Line 3. But since task $t_2$ writes to `y` *after* signaling, then the read from task $t_1$ in Line 12 runs concurrently with the write of task $t_2$ in Line 4.

A feature that distinguishes Habanero phasers from other barrier-like synchronization mechanisms is that waiting for signalers is *optional*. A task can choose to manipulate a phaser according to two abilities: (i) the ability to observe phaser synchronization, *i.e.,* waiter, and (ii) the ability to influence synchronization, *i.e.,* signaler. Each member is registered according to a mode $r$ among: `SW` for tasks that must signal and wait, `WO` for tasks that wait but do not signal, and `SO` for tasks that signal but do not wait. In the example, task $t_1$ registers task $t_2$ in phaser `ph` using mode `SW`, which is short for `SIG_WAIT`, given by expression `ph.inMode(SIG_WAIT)` in Line 2.

Waiting observes the signals from *every* signaler. Thus, tasks that wait and signal, mode `SW`, as in the example, must signal before waiting at every phase to prevent waiting for a signal the task did not produce. This synchronization pattern is known as a *barrier*. Members disregard wait-only (`WO`) tasks upon waiting; this subset of tasks cannot influence synchronization, only observe it. Phasers can encode

$$\frac{P(t)=v \qquad \text{Signaler } v \qquad v.\texttt{mode}=\texttt{SW} \implies v.\texttt{wp}=v.\texttt{sp}}{P \xrightarrow{t:\texttt{signal}} P[t \mapsto v.\texttt{sp} := v.\texttt{sp}+1]}$$

$$\frac{t \text{ sync } P \qquad P(t)=v \qquad \text{Waiter } v \qquad v.\texttt{mode}=\texttt{SW} \implies v.\texttt{wp}+1=v.\texttt{sp}}{P \xrightarrow{t:\texttt{wait}} P[t \mapsto v.\texttt{wp} := v.\texttt{wp}+1]}$$

$$\frac{t' \notin P \qquad P(t)=v \qquad \text{Waiter } r \implies \text{Waiter } v \qquad \text{Signaler } r \implies \text{Signaler } v}{P \xrightarrow{t:\texttt{reg}(t',r)} P[t' \mapsto v.\texttt{mode} := r]}$$

$$\frac{t \in P}{P \xrightarrow{t:\texttt{drop}} P - t}$$

Figure 2: Operational semantics of phaser operations

latches, future-promises, and fork-join synchronization patterns using wait-only tasks. Finally, tasks that only signal, do not wait for others; this lets phasers encode producer-consumer synchronization.

**HJ Phaser formalization.** We define the state of a phaser $P$ to be a map from members $\mathscr{T}$ into *views* $\mathscr{V}$, which holds the signal count, the wait count, and registration mode of a member. Consider the usual operations on finite maps (which we use to encode phasers) with the given notation: predicate $P(t)=v$ ensures that the pair of key $t$ and value $v$ is a member of map $P$, predicate $t \in P$ is short-hand for $\exists v\colon P(t)=v$, map $P[t \mapsto v]$ adds the pair $t$ and $v$ to map $P$ (replacing the assigned view if $t \in P$), and map $P-t$ results from removing the pair associated with key $t$ from map $P$. In the mechanization, we use Coq's standard library of finite maps Coq.FSets.FMaps.

A view $v$ represents the task-local information that each member has over the phaser. The view consists of a triple: the first value $n$ is a natural number that counts the number of times the given task issued a signal on the target phaser and can be accessed by $v.\texttt{sp}$; the second value $m$ counts the number of waits and can be accessed by $v.\texttt{wp}$; the third value $r$ is the registration mode of the given task and is accessed by $v.\texttt{mode}$.

$$v ::= \{\texttt{sp} := n, \texttt{wp} := m, \texttt{mode} := r\}$$

The field update operation $v.f := e$ yields a view that is the same as $v$ except for field $f$ that becomes $e$. For instance, $v.\texttt{sp} := 3$ yields a view, say $w$, where $w.\texttt{wp} = v.\texttt{wp}$, $w.\texttt{sp} = 3$, and $w.\texttt{mode} = v.\texttt{mode}$. To inquire the signaling and waiting abilities of a view we have the following predicates: Waiter $r \stackrel{\text{def}}{=} r \in \{\texttt{WO}, \texttt{SW}\}$, Signaler $r \stackrel{\text{def}}{=} r \in \{\texttt{SO}, \texttt{SW}\}$. And let the short-hand notation Signaler $v \stackrel{\text{def}}{=}$ Signaler $v.\texttt{mode}$ and Waiter $v \stackrel{\text{def}}{=}$ Waiter $v.\texttt{mode}$.

We define a small-step operational semantics for phaser operations in Figure 2. The reduction $\xrightarrow{t:o}$ is labeled by the member $t$ issuing the operation, and by an operation $o$ defined below.

$$o ::= \texttt{signal} \mid \texttt{wait} \mid \texttt{reg}(t,r) \mid \texttt{drop}$$

**Remark 1.** *To model Java phasers and X10 clocks semantics refer to Figure 2 but limit the registration mode to signal-wait mode, that is* $r ::= \texttt{SW}$.

Operation `signal` increments the signal phase. Only tasks registered as signalers can issue this operation. The pre-conditions in signal, $v.\texttt{wp} = v.\texttt{sp}$, and in wait, $v.\texttt{wp}+1 = v.\texttt{sp}$, enforce tasks registered in signal-wait mode to interleave each signal with a wait.

Waiting is the crux of synchronization; this is captured by $t$ sync $P$, defined next.

$$\frac{P(t).\texttt{mode} = \texttt{SO}}{t \text{ sync } P} \qquad \frac{\text{Waiter } P(t) \quad \text{Await}(P, P(t).\texttt{wp}+1)}{t \text{ sync } P}$$

Signal-only tasks do not wait for others, so $t$ sync $P$ holds in this case. Waiter task $t$ must await the subsequent wait-phase $P(t).\texttt{wp}+1$. Proposition $\text{Await}(P,n)$ holds once phase $n$ can be observed; the definition ensures that all tasks that can signal have issued at least $n$ signals.

$$\text{Await}(P,n) \stackrel{\text{def}}{=} \forall t: \ \text{Signaler } P(t) \implies P(t).\texttt{sp} \geq n$$

Tasks register other tasks with with `asyncPhased`, which is captured by $\texttt{reg}(t,r)$. For instance, instruction `asyncPhased(ph.inMode(SIG_WAIT),...)` in Line 1 becomes $\texttt{reg}(t_2, \texttt{SW})$ in this semantics if we are spawning task $t_2$. Habanero phasers limit task registration: only unregistered tasks can be added, thus $t' \in P$, only registered tasks $t$ can add new members, $P(t) = v$, only waiters can register other waiters, hence Waiter $r \implies$ Waiter $v$, and only signalers can register other signalers, so Signaler $r \implies$ Signaler $v$.

To establish the results in the next section, let us establish the invariant of well-formedness. Additionally, let $P \to Q$ be defined as there exist $t$ and $o$ such that $P \xrightarrow{t:o} Q$.

**Definition 1** (Well-formed view)**.** *Let a well-formed view be such that $v.\texttt{wp} \leq v.\texttt{sp}$ and if* Waiter $v$ *then $v.\texttt{sp} - v.\texttt{wp} \leq 1$. Let $\mathscr{V}^{WF}$ be the set of all well-formed views.*

**Lemma 1** (Reduction preserves well-formedness of views)**.** *Let $P$ be such that if $P(t) = v$, then $v \in \mathscr{V}^{WF}$. If $P \to Q$, then $Q$ is such that if $Q(t) = v$, then $v \in \mathscr{V}^{WF}$.*

Henceforth, we only consider views that are in $\mathscr{V}^{WF}$.

## 3 Phase Ordering

This section formalizes Phase Ordering, originally introduced in [13], to reason about whether two tasks should execute concurrently in terms of views $\mathscr{V}$ and states $\mathscr{P}$. Specifically, Phaser Ordering is a Happens-Before relation: if the number of signals issued by a task is smaller than the last phase observed by some other task, then the former Happened Before the latter. For instance, let $v_1$ be a view that task $t_1$ has over the phaser in Figure 1 and $v_2$ be a view that task $t_2$ has over the phaser in Figure 1, each from a distinct state of the same phaser `ph`, *i.e.,* there exists two states $P$ and $Q$ such that $P(t_1) = v_1$ and $Q(t_2) = v_2$. Now, let $v_1 \stackrel{\text{def}}{=} \{\texttt{sp} := 0, \texttt{wp} := 0, \texttt{mode} := \texttt{SW}\}$ be the view of $t_1$ when executing `x:=1` and $v_2 \stackrel{\text{def}}{=} \{\texttt{sp} := 1, \texttt{wp} := 1, \texttt{mode} := \texttt{SW}\}$ be the view $t_2$ when executing `println(y)`. The registration mode tells us that view $v_2$ must observe and wait for the signals of $v_1$. View $v_1$ tells us that $t_1$ did not produce any signal and $v_2$ tells us that $t_2$ observed phase 1 (a collective signal). Thus, since $t_1$ signaled fewer times than the phase observed by $t_2$, we can infer that $v_1$ must have happened before $v_2$. In order for view $v_2$ to observe a signal, the task controlling view $v_1$ must eventually signal to become $v_1.\texttt{sp} = 1$.

**Definition 2** (Happens-before (HB) relation)**.** *Let $v_1 \prec v_2$ read as $v_1$ must have happened before $v_2$, defined as the conjunction of:*

$$\text{Signaler } v_1 \qquad v_1.\texttt{sp} < v_2.\texttt{wp} \qquad \text{Waiter } v_2$$

*We say that $P \prec Q$ if there exist two tasks $t, t'$ such that $P(t) \prec Q(t')$.*

**Example 1.** *Suppose* $P \xrightarrow{t:\texttt{signal}} Q \xrightarrow{t:\texttt{wait}} R$ *and that* $P(t).\mathit{mode} = \texttt{SW}$. *We have that* $P \prec R$.

*Proof.* Let $P(t) = v$, $Q(t) = w$, and $R(t) = u$. First, we simplify our goal, since we know that from `wait` $u.\texttt{wp} = w.\texttt{wp} + 1$ and because `signal` does not alter the wait phase we have that $w.\texttt{wp} = v.\texttt{wp}$. Hence, $v.\texttt{sp} < u.\texttt{wp} \equiv v.\texttt{sp} < w.\texttt{wp} + 1 \equiv v.\texttt{sp} < v.\texttt{wp} + 1$. Now, by inverting reduction $\xrightarrow{t:\texttt{signal}}$, we get two cases: either $v.\texttt{wp} = v.\texttt{sp}$ and it trivially holds, otherwise we get a contradiction. □

Let us show that $\prec$ is a causality relation over views, a fundamental notion for many problems occurring in distributed computing [11]. By causality relation we mean a strict partial order: (i) *transitive:* if $v_1 \prec v_2$ and $v_2 \prec v_3$, then $v_1 \prec v_3$; (ii) *irreflexive:* for all $v$, we have that $\neg(v \prec v)$; (iii) *asymmetric:* if $v_1 \prec v_2$, then $\neg(v_2 \prec v_1)$.

**Lemma 2.** $(\prec, \mathscr{V})$ *is a causality relation.*

To show that $\prec$ is a causality relation over phasers, we need to establish some auxiliary results that reason about states that cannot happen before others, $\neg(P \prec Q)$. Since Coq uses a constructive logic, it is easier avoid the use of false in our premises. Let the negation of Happens-Before be defined as Cannot-Happen-Before.

**Definition 3** (Cannot-Happen-Before relation)**.** *Let* $v_1 \trianglerighteq v_2$ *read as* $v_1$ *cannot happen before* $v_2$*:*

$$v_1.\mathit{mode} = \texttt{WO} \quad \vee \quad v_1.\mathit{sp} \geq v_2.\mathit{wp} \quad \vee \quad v_2.\mathit{mode} = \texttt{SO}$$

*We say that* $P \trianglerighteq Q$ *if for any tasks* $t, t'$ *we have that* $P(t) \trianglerighteq Q(t')$.

Although our proofs use $\trianglerighteq$, the following remark allow us to present our lemmas with the more familiar HB relation, $\neg(P \prec Q)$.

**Remark 2.** *We have that* $v_1 \prec v_2 \iff \neg(v_1 \trianglerighteq v_2)$, $v_1 \trianglerighteq v_2 \iff \neg(v_1 \prec v_2)$, $P_1 \prec P_2 \implies \neg(P_1 \trianglerighteq P_2)$, *and* $P_1 \trianglerighteq P_2 \implies \neg(P_1 \prec P_2)$

Finally, we define the usual notion of May-Happen-in-Parallel (or concurrency relation) for views and for phasers.

**Definition 4** (May-Happen-in-Parallel relation)**.** *Let* $v_1 \,||\, v_2$ *read as* $v_1$ *happens in parallel with* $v_2$ *and be defined as* $v_1 \trianglerighteq v_2$ *and* $v_2 \trianglerighteq v_1$. *Let* $P_1 \,||\, P_2$ *be defined as* $P_1 \trianglerighteq P_2$ *and* $P_2 \trianglerighteq P_1$.

# 4 Results

Similarly to what happened with showing the causality of $\prec$ over views, when establishing the causality of $\prec$ over phasers we require an invariant that relates the various local views of a phaser. While in the context of views, we must ensure that the wait phase does not overtake the signal phase, in the context of phasers we must ensure that its views may happen in parallel with each other, *i.e.,* there must be no scheduling constraints within a phaser.

**Definition 5** (Well-ordered phaser)**.** *Let a* well-ordered phaser *be such that* $P \,||\, P$. *Let* $\mathscr{P}^{WO}$ *be the set of all well-ordered phasers.*

Reduction *cannot* introduce unsolvable scheduling constraints within a phaser.

**Lemma 3** (Reduction preserves phaser well-orderedness)**.** *If* $P \in \mathscr{P}^{WO}$ *and* $P \to Q$, *then* $Q \in \mathscr{P}^{WO}$.

We are now ready to show that $\prec$ is a causality relation over phasers.

**Theorem 1.** $(\prec, \mathscr{P}^{WO})$ *is a causality relation.*

The execution of an HJ program must respect the scheduling restriction imposed by Phase Ordering. In the point of view of our formalization, the reduction relation captures the execution of a single phaser operation. Thus, Theorem 2 shows that the state after execution cannot happen before the state before execution, or, in other words, that the pre- and post-states of a phaser operation respect Phase Ordering.

**Theorem 2.** *If* $P \in \mathscr{P}^{WO}$ *and* $P \to Q$ *we have that* $\neg(Q \prec P)$.

It is curious to consider that from $P \to Q$ we can also conclude $\neg(P \prec Q)$. Thus, it follows that.

**Lemma 4.** *If* $P \in \mathscr{P}^{WO}$, $P \to Q$, *then* $P \,||\, Q$.

Be aware, however, that MHP does *not* enjoy transitivity, otherwise HB would be an empty relation! Example 1 is an evidence of when $P \,||\, Q$ and $Q \,||\, R$ but $\neg(P \,||\, R)$, as $P \prec R$. This also tells us that for any two states $P$ and $Q$ if we have $P \prec Q$, then state $Q$ is the result of at least two phaser operations.

The final theorem establishes that the execution of an HJ program respects the scheduling restriction imposed by Phase Ordering. While Theorem 2 relates the pre- and post-states of executing a single operation, Theorem 3 generalizes this result to any possible execution trace, showing that Phase Ordering captures the execution order of instructions in programs that use phasers. Let $\to^*$ be defined as the reflexive transitive closure of $\to$.

**Theorem 3** (Absence of synchronization errors)**.** $P \in \mathscr{P}^{WO}$, $P \to^* Q$, *then* $\neg(Q \prec P)$.

Our results can be summarized into three groups. The first group consists of Lemmas 1 and 3; this serves as a steppingstone for our main results. The former lemma establishes an invariant on a relationship between wait and signal phases ($\mathscr{V}^{WF}$), while the latter lemma establishes an invariant on a relationship between any two views picked from a state ($\mathscr{P}^{WO}$). Since $\mathscr{V}^{WF}$ and $\mathscr{P}^{WO}$ are preserved by our semantics, any program that manipulates Habanero phasers can assume Lemmas 1 and 3 to hold. The second group consists of Lemma 2 and Theorem 1 and lets us relate views (and states) with the HB relation. The third group consists of Theorems 2 and 3 and it lets us conclude that the execution of a program using Habanero phasers respects the scheduling restriction imposed by Phase Ordering (HB) relation.

## 5 Conclusion

In this paper we propose the first formalization of Habanero phaser semantics and of Phase Ordering, from which we derive the May-Happen-In-Parallel (MHP) and Happens-Before (HB) relations for phaser operations, as part of an ongoing effort to formalize the Habanero programming model. Our definitions and proofs are mechanized using the Coq proof assistant, it consists of 2600 lines of code and 140 lemmas. Our next step is to verify data-race errors in parallel programs that feature collective producer-consumer synchronization patterns.

## Acknowledgments

We thank Nick Vrvilo and the anonymous reviewers for their comments and suggestions.

# References


[1] Yan Cai, Shangru Wu & W. K. Chan (2014): *ConLock: A Constraint-based Approach to Dynamic Checking on Deadlocks in Multithreaded Programs*. In: *ICSE'14*, ACM, pp. 491–502, doi:10.1145/2568225.2568312.

[2] Philippe Charles, Christian Grothoff, Vijay Saraswat, Christopher Donawa, Allan Kielstra, Kemal Ebcioglu, Christoph von Praun & Vivek Sarkar (2005): *X10: an object-oriented approach to non-uniform cluster computing*. In: *OOPSLA'05*, ACM, pp. 519–538, doi:10.1145/1103845.1094852.

[3] Tiago Cogumbreiro (2016): *Habanero Coq formalization project*. Available at `https://github.com/cogumbreiro/habanero-coq/tree/places16`.

[4] Tiago Cogumbreiro, Raymond Hu, Francisco Martins & Nobuko Yoshida (2015): *Dynamic Deadlock Verification for General Barrier Synchronisation*. In: *PPoPP'15*, ACM, pp. 150–160, doi:10.1145/2688500.2688519.

[5] Silvia Crafa, David Cunningham, Vijay Saraswat, Avraham Shinnar & Olivier Tardieu (2014): *Semantics of (Resilient) X10*. In Richard Jones, editor: *ECOOP'14*, *LNCS* 8586, Springer, pp. 670–696, doi:10.1007/978-3-662-44202-9_27.

[6] Evelyn Duesterwald & Mary Lou Soffa (1991): *Concurrency Analysis in the Presence of Procedures Using a Data-flow Framework*. In: *TAV'91*, ACM, pp. 36–48, doi:10.1145/120807.120811.

[7] Tayfun Elmas, Shaz Qadeer & Serdar Tasiran (2006): *Goldilocks: Efficiently Computing the Happens-before Relation Using Locksets*. In: *FATES'06/RV'06*, Springer, pp. 193–208, doi:10.1007/11940197_13.

[8] Saurabh Joshi, Rudrapatna K. Shyamasundar & Sanjeev K. Aggarwal (2012): *A New Method of MHP Analysis for Languages with Dynamic Barriers*. In: *IPDPSW'12*, IEEE, pp. 519–528, doi:10.1109/IPDPSW.2012.70.

[9] Leslie Lamport (1978): *Time, Clocks, and the Ordering of Events in a Distributed System*. *Communications of the ACM* 21(7), pp. 558–565, doi:10.1145/359545.359563.

[10] Pallavi Maiya, Aditya Kanade & Rupak Majumdar (2014): *Race Detection for Android Applications*. In: *PLDI '14*, ACM, pp. 316–325, doi:10.1145/2594291.2594311.

[11] Reinhard Schwarz & Friedemann Mattern (1994): *Detecting Causal Relationships in Distributed Computations: In Search of the Holy Grail*. *Distributed Computing* 7(3), pp. 149–174, doi:10.1007/BF02277859.

[12] Rahul Sharma, Michael Bauer & Alex Aiken (2015): *Verification of Producer-consumer Synchronization in GPU Programs*. In: *PLDI'15*, ACM, pp. 88–98, doi:10.1145/2737924.2737962.

[13] Jun Shirako, David M. Peixotto, Vivek Sarkar & William N. Scherer (2008): *Phasers: a unified deadlock-free construct for collective and point-to-point synchronization*. In: *ICS'08*, ACM, pp. 277–288, doi:10.1145/1375527.1375568.

[14] Yannis Smaragdakis, Jacob Evans, Caitlin Sadowski, Jaeheon Yi & Cormac Flanagan (2012): *Sound Predictive Race Detection in Polynomial Time*. In: *POPL'12*, ACM, pp. 387–400, doi:10.1145/2103656.2103702.

[15] Tomofumi Yuki, Paul Feautrier, Sanjay V. Rajopadhye & Vijay Saraswat (2013): *Checking Race Freedom of Clocked X10 Programs*. *CoRR* abs/1311.4305. Available at `http://arxiv.org/abs/1311.4305`.

# A Main definitions and results in Coq

This section presents code listings, in Coq syntax, of the definitions and results found in Sections 3 and 4. The full proof scripts and auxiliary lemmas can be found online in our open source project HJ-Coq [3]. The next code listing shows definitions related to views.

```
(* Module declares a name space to avoid avoid name collisions
   between the definitions of views and phasers. *)
Module Taskview.
 (* Definition 1 for views, defined by three cases: *)
 Inductive Wellformed v : Prop :=
   | tv_wellformed_wait_cap_eq:
     WaitCap (mode v) →
     wait_phase v = signal_phase v →
     Wellformed v
   | tv_wellformed_wait_cap_succ:
     WaitCap (mode v) →
     S (wait_phase v) = signal_phase v →
     Wellformed v
   | tv_wellformed_so:
     mode v = SIGNAL_ONLY →
     wait_phase v <= signal_phase v →
     Wellformed v.
 (* Definition 2 for views: *)
 Inductive HappensBefore v1 v2 : Prop :=
   tv_hb_def:
     signal_phase v1 < wait_phase v2 →
     SignalCap (mode v1) →
     WaitCap (mode v2) →
     HappensBefore v1 v2.
 (* Declare the infix notation of the HappensBefore relation *)
 Infix "≺" := HappensBefore : phaser_scope.
 (* Definition 3 for views, defined by three cases: *)
 Inductive CannotHappenBefore v1 v2 : Prop :=
   | tv_chb_ge:
     signal_phase v1 >= wait_phase v2 →
     CannotHappenBefore v1 v2
   | tv_chb_so:
     mode v2 = SIGNAL_ONLY →
     CannotHappenBefore v1 v2
   | tv_chb_wo:
     mode v1 = WAIT_ONLY →
     CannotHappenBefore v1 v2.
 (* Define the infix notation of HappensBefore. *)
 Infix "⊵" := CannotHappenBefore : phaser_scope.
End Taskview.
```

Next, we list the definitions and notations related to phasers from Section 3.

```
Module Phaser.
  (* Definition 1 for phasers: *)
  Inductive Wellformed (ph:phaser) : Prop :=
    ph_wellformed_def:
```

```
      (∀ t v, Map_TID.MapsTo t v ph → Taskview.Wellformed v) →
      Wellformed ph.
  (* Definition 2 for phasers: *)
  Inductive HappensBefore (ph1 ph2:phaser) : Prop :=
    ph_hb_def:
      ∀ t1 t2 v1 v2,
      Map_TID.MapsTo t1 v1 ph1 →
      Map_TID.MapsTo t2 v2 ph2 →
      Taskview.HappensBefore v1 v2 →
      HappensBefore ph1 ph2.
  (* Defines the infix notation of HappensBefore. *)
  Infix "≺" := HappensBefore : phaser_scope.
  (* Definition 3 for phasers: *)
  Inductive CannotHappenBefore (ph1 ph2:phaser) : Prop :=
    ph_chb_def:
      (∀ t1 t2 v1 v2, Map_TID.MapsTo t1 v1 ph1 → Map_TID.MapsTo t2 v2 ph2 →
        Taskview.CannotHappenBefore v1 v2) →
      CannotHappenBefore ph1 ph2.
  (* Defines the infix notation of CannotHappenBefore. *)
  Infix "⊵" := CannotHappenBefore : phaser_scope.
  (* Definition 5: *)
  Inductive WellOrdered x : Prop :=
    well_ordered_def: Facilitates x x → WellOrdered x.
  (* Definition 5: *)
  Inductive Par x y: Prop :=
    par_def: Facilitates x y → Facilitates y x → Par x y.
End Phaser.
```

Finally, we cross-reference the lemmas and theorems in the paper against the Coq mechanization.

```
(* Lemma 1 *)
Lemma ph_reduces_preserves_wellformed:
∀ ph t o ph', Wellformed ph → Reduces ph t o ph' → Wellformed ph'.
(* Lemma 2 *)
Theorem tv_lt_trans: ∀ x y z, Wellformed y → x ≺ y → y ≺ z → x ≺ z.
Theorem tv_lt_antisym:∀ x y, Wellformed x → Wellformed y → x ≺ y → ¬(y ≺ x).
Theorem tv_lt_irreflexive: ∀ v, Wellformed v → ¬(v ≺ v).
(* Lemma 3 *)
Lemma ph_hb_irreflexive: ∀ ph, WellOrdered ph → ¬(ph cc≺ ph).
Lemma ph_hb_antisym: ∀ x y, WellOrdered x → WellOrdered y → x ≺ y → ¬(y ≺ x).
Lemma ph_hb_trans: ∀ ph1 ph2 ph3, WellOrdered ph2 → ph1 ≺ ph2 → ph2 ≺ ph3 → ph1 ≺ ph3
(* Theorem 1 *)
Lemma reduces_ne: ∀ x y, WellOrdered x → SReduces x y → x ⊵y.
(* Lemma 4 *)
Lemma reduces_par: ∀ x y, Wellformed x → WellOrdered x → SReduces x y → x || y.
(* Theorem 2 *)
Lemma ph_ge_reduce:
∀ ph t o ph', Wellformed ph → WellOrdered ph → Reduces ph t o ph' → ph' ⊵ph.
(* Theorem 3 *)
Lemma ph_s_reduces_trans_refl_ge: ∀ x y, Wellformed x → WellOrdered x →
clos_refl_trans phaser SReduces x y → y ⊵x.
```

# B Proof sketches

The proofs for all lemmas and theorems in this paper are machine checked in [3]. In this section, we show the proof sketches for the main results.

**Theorem 1** $(\prec, \mathscr{P}^{WO})$ is a causality relation.

*Proof.* $(\prec, \mathscr{P}^{WO})$ is transitive: if $P \prec Q$ and $Q \prec R$, then $P \prec R$. We invert $P \prec Q$ and get that there exists $t_1$ and $t_2$ such that $P(t_1) = v_1$, $Q(t_2) = v_2$, and $v_1 \prec v_2$. Similarly, we invert $Q \prec R$ and get that there exists $t_3$ and $t_4$ such that $Q(t_3) = v_3$, $R(t_4) = v_4$ such that $v_3 \prec v_4$. From $Q \in \mathscr{P}^{WO}$, $Q(t_3) = v_3$, and $Q(t_2) = v_2$, thus $v_3 \trianglerighteq v_2$. From $v_1 \prec v_2$, $v_3 \prec v_4$, and $v_3 \trianglerighteq v_2$, we can conclude that $v_1 \prec v_4$, and therefore $P \prec R$.

$(\prec, \mathscr{P}^{WO})$ is irreflexive: if $P \in \mathscr{P}^{WO}$ then $\neg(P \prec P)$. From $P \in \mathscr{P}^{WO}$ we get that $P \trianglerighteq P$. Then, we apply Remark 2 and get that $\neg(P \prec P)$.

$(\prec, \mathscr{P}^{WO})$ is asymmetric: if $P \prec Q$, then $\neg(Q \prec P)$. Using Remark 2 it is enough to show that $Q \trianglerighteq P$, specifically that if $Q(t_1) = v_1$ and $P(t_2) = v_2$, then $v_1 \trianglerighteq v_2$. It can be shown that for any pair of views we have that $v_1 \prec v_2$ or $v_1 \trianglerighteq v_2$. Since the latter concludes the proof directly, we proceed to show that the former case, $v_1 \prec v_2$, leads to a contradiction. From $P \prec Q$ we have that there exists task $t$ and $t'$ such that $P(t) = v$, $Q(t') = w$, and $v \prec w$. The contradiction arises from arriving at $\neg(v_1 \trianglerighteq w)$ and $v_1 \trianglerighteq w$. First, we get $v_1 \trianglerighteq w$ by applying Definition 5 to $P \in \mathscr{P}^{WO}$, $Q(t_1) = v_1$, and $Q(t') = w$. Second, we show $\neg(v_1 \trianglerighteq w)$. Applying Definition 5 to $P \in \mathscr{P}^{WO}$, $P(t) = v$, and $P(t_2) = v_2$ results in $v \trianglerighteq v_2$. As we have seen in the proof of transitivity, from $v_1 \prec v_2$, $v \prec w$, and $v \trianglerighteq v_2$, we conclude $v_1 \prec w$. Applying Remark 2 we get that $\neg(v_1 \trianglerighteq w)$, which leads to the contradiction. □

**Theorem 2** If $P \in \mathscr{P}^{WO}$ and $P \to Q$ we have that $\neg(Q \prec P)$.

*Proof.* By inverting the hypothesis $P \to Q$ we get $P \xrightarrow{t:o} Q$. Next, we invert $P \xrightarrow{t:o} Q$ and obtain four cases, one for each constructor of $o$. For each case it is enough to show that given $P(x) = v_x$ and $Q(y) = v_y$, we can obtain $v_y \trianglerighteq v_x$.

Cases $o = \texttt{signal}$ and $o = \texttt{wait}$ proceed similarly. We test if $y = t$. If $y \neq t$, then $Q(y) = P(y)$. We can obtain $v_y \trianglerighteq v_x$ from $Q(y) = v_y$ , $Q(y) = v_y$ and $Q \in \mathscr{P}^{WO}$ (which we get from Lemma 3, $P \to Q$, and $P \in \mathscr{P}^{WO}$). Otherwise $y = t$, and we get that there exists a view $v$ such that $P(y) = v$. Let the increment of the signal phase (wait phase) be denoted by $o(v)$ where $o(v) = v_y$. At this point, we have $P(y) = v$ and $Q(y) = o(v)$. Then, we just need to show that $o(v_x) \trianglerighteq v_x$, given that $v \trianglerighteq v_x$, which we get from $P(y) = v$, $P(x) = v_x$, and $P \in \mathscr{P}^{WO}$.

Case $o = \texttt{reg}(t', r)$, where we have $P(t) = v$. We test if $t' = y$. If $t' \neq y$, then $P(y) = v_y$. Thus, we have $v_y \trianglerighteq v_x$ from $P(y) = v_y$, $P(x) = v_x$, and $P \in \mathscr{P}^{WO}$. Otherwise, $t' = y$, and therefore $v_y = (v\texttt{.mode :=}\, r)$. From $P(t) = v$, $P(x) = v_x$, and $P \in \mathscr{P}^{WO}$, we conclude $v \trianglerighteq v_x$. We close the case by showing that from $v \trianglerighteq v_x$ then $(v\texttt{.mode :=}\, r) \trianglerighteq v_x$ holds.

Case $o = \texttt{drop}$, since we have $Q = P - t$, then $P(y) = v_y$. From $P(y) = v_y$, $P(x) = v_x$, and $P \in \mathscr{P}^{WO}$, we get $v_y \trianglerighteq v_x$. □

**Theorem 3** $P \in \mathscr{P}^{WO}$, $P \to^* Q$, then $\neg(Q \prec P)$.

*Proof.* We state our result in terms of $Q \trianglerighteq P$ and, but we can use Remark 2 to obtain $\neg(Q \prec P)$. The proof follows by induction on the derivation tree of $P \rightarrow^* Q$. For the base case we have that $Q = P$, and we show that $P \trianglerighteq P$ holds.

For the inductive case we have that there exists a phaser $R$ such that $P \rightarrow^* R$, $R \rightarrow Q$, and $R \trianglerighteq P$. From $P \in \mathscr{P}^{WF}$ (which states that every view $v$ in $R$ is $v \in \mathscr{V}^{WF}$) and $P \rightarrow^* R$ we can get that $R \in \mathscr{P}^{WF}$, by performing induction on the structure of $P \rightarrow^* R$ and using Lemma 1. Similarly, we get that $R \in \mathscr{P}^{WO}$, by performing induction on the structure of $P \rightarrow^* R$ and using Lemma 2. The final step of the proof is to show that if $R \in \mathscr{P}^{WF}$ $R \trianglerighteq P$, and $R \rightarrow Q$, then $Q \trianglerighteq P$. At this point, we perform a case analysis in the reduction relation $R \rightarrow Q$. The case for $R \xrightarrow{t:\texttt{drop}} Q$ is trivial, thus we shift our attention to the remaining cases, where dom$R \subseteq$ dom$Q$, and the crux of the proof is showing that if $P(x) = v_x$ and $R(y) = v_y$, and $Q(z) = v_z$, then $v_z \trianglerighteq v_x$.

We now check if $t = z$. If the check succeeds, then we can conclude that there is a view $v$ such that $R(z) = v$ such that $v_x$ results by applying a signal or wait to $v$, notation $v_x = o(v)$. Now, because we have that $R \trianglerighteq P$, then $v \trianglerighteq v_x$, so we just need to show that $o(v) \trianglerighteq v_x$, which we omit detailing.

Finally, we address the case where $t \neq z$. We inspect $o$ and discuss the non-trivial case, when there exist $t'$, $r$, and $v$ such that $o = \texttt{reg}(t', r)$, $R(t) = v$, and $v_z = (v.\texttt{mode} := r)$. Recall, we want to show that $(v.\texttt{mode} := r) \trianglerighteq v_x$ holds. The proof can be concluded using three premises: Waiter $r \implies$ Waiter $v$ and Signaler $r \implies$ Signaler $v$, which we get from the reduction rule on $\texttt{reg}(t', r)$; and $v \trianglerighteq v_x$, which we get from $R \trianglerighteq P$. □